

FORM PTO-1390 (REV 11-98)	U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE		ATTORNEY'S DOCKET NUMBER 100-43
TRANSMITTAL LETTER TO THE UNITED STATES DESIGNATED/ELECTED OFFICE (DO/EO/US) CONCERNING A FILING UNDER 35 U.S.C. 371		U.S. APPLICATION NO. (If known, see 37 C.F.R. 1.5)	
INTERNATIONAL APPLICATION NO. PCT/AU99/00297		INTERNATIONAL FILING DATE 20 April 1999	PRIORITY DATE CLAIMED 20 April 1998
TITLE OF INVENTION ACCESS CONTROL METHOD AND APPARATUS			
APPLICANT(S) FOR DO/EO/US Michael Peter RUMSEWICZ and Maurice David CASTRO			
Applicant herewith submits to the United States Designated/Elected Office (DO/EO/US) the following items and other information:			
<ol style="list-style-type: none"> 1. <input checked="" type="checkbox"/> This is a FIRST submission of items concerning a filing under 35 U.S.C. 371. 2. <input type="checkbox"/> This is a SECOND or SUBSEQUENT submission of items concerning a filing under 35 U.S.C. 371. 3. <input checked="" type="checkbox"/> This is an express request to begin national examination procedures (35 U.S.C. 371(f) at any time rather than delay examination until the expiration of the applicable time limit set in 35 U.S.C. 371(b) and PCT Articles 22 and 39(1). 4. <input checked="" type="checkbox"/> A proper Demand for International Preliminary Examination was made by the 19th month from the earliest claimed priority date. 5. A copy of the International Application as filed (35 U.S.C. 371(c)(2)). <ol style="list-style-type: none"> a. <input type="checkbox"/> is transmitted herewith (required only if not transmitted by the International Bureau). b. <input checked="" type="checkbox"/> has been transmitted by the International Bureau. c. <input type="checkbox"/> is not required, as the application was filed in the United States Receiving Office (RO/US). 6. <input type="checkbox"/> A translation of the International Application into English (35 U.S.C. 371(c)(2)). 7. <input type="checkbox"/> Amendments to the claims of the International Application under PCT Article 19 (35 U.S.C. 371(c)(3)). <ol style="list-style-type: none"> a. <input type="checkbox"/> are transmitted herewith (required only if not transmitted by the International Bureau). b. <input type="checkbox"/> have been transmitted by the International Bureau. c. <input type="checkbox"/> have not been made; however, the time limit for making such amendments has NOT expired. d. <input type="checkbox"/> have not been made and will not be made. 8. <input type="checkbox"/> A translation of the amendments to the claims under PCT Article 19 (U.S.C. 371(c)(3)). 9. <input type="checkbox"/> An oath or declaration of the inventor(s) (35 U.S.C. 371(c)(4)). 10. <input type="checkbox"/> A translation of the annexes to the International Preliminary Examination Report under PCT Article 36 (35 U.S.C. 371(c)(5)). 			
Items 11. To 16. Below concern document(s) or information included: <ol style="list-style-type: none"> 11. <input type="checkbox"/> An Information Disclosure Statement under 37 C.F.R. 1.97 and 1.98. 12. <input type="checkbox"/> An assignment document for recording. A separate cover sheet in compliance with 37 C.F.R. 3.28 and 3.31 is included. 13. <input checked="" type="checkbox"/> A FIRST preliminary amendment. <input type="checkbox"/> A SECOND or SUBSEQUENT preliminary amendment. 14. <input type="checkbox"/> A substitute specification. 15. <input type="checkbox"/> A change of power of attorney and/or address letter. 16. <input checked="" type="checkbox"/> Other items or information. International Search Report 			

U.S. APPLICATION NUMBER (if known, see 37 C.F.R. 1.6)	INTERNATIONAL APPLICATION NO. PCT/AU99/00297	ATTORNEY'S DOCKET NUMBER 100-43			
09/1673609 to be assigned		CALCULATIONS PTO USE ONLY			
<p>17. <input checked="" type="checkbox"/> The following fees are submitted:</p> <p>BASIC NATIONAL FEE (37 C.F.R. 1.492(a)(1)-(5):</p> <ul style="list-style-type: none"> -- Neither international preliminary examination fee (37 C.F.R. 1.482) nor international search fee (37 C.F.R. 1.445(a)(2)) paid to USPTO and International Search Report not prepared by the EPO or JPO.....\$1000.00 -- International preliminary examination fee (37 C.F.R. 1.482) not paid to USPTO but International Search Report prepared by the EPO or JPO\$860.00 -- International preliminary examination fee (37 C.F.R. 1.482) not paid to USPTO but international search fee (37 C.F.R. 1.445(a)(2)) paid to USPTO.....\$710.00 -- International preliminary examination fee paid to USPTO (37 C.F.R. 1.482) but all claims did not satisfy provisions of PCT Article 33(1)-(4)\$690.00 -- International preliminary examination fee paid to USPTO (37 C.F.R. 1.482) and all claims satisfied provisions of PCT Article 33(1)-(4)\$100.00 					
ENTER APPROPRIATE BASIC FEE AMOUNT = <input type="text" value="1000.00"/>					
<p>Surcharge of \$130.00 for furnishing the oath or declaration later than <input type="checkbox"/> 20 <input checked="" type="checkbox"/> 30 months from the earliest claimed priority date (37 C.F.R. 1.492(e)).</p>					
<input type="text" value="\$ 130.00"/>					
CLAIMS	NUMBER FILED	NUMBER EXTRA			
Total Claims	46	20 = 26			
Independent Claims	2	-3 = 0			
MULTIPLE DEPENDENT CLAIMS(S) (if applicable) <input type="text" value="\$270.00"/>					
TOTAL OF ABOVE CALCULATIONS = <input type="text" value="\$ 1598.00"/>					
<p>Reduction by ½ for filing by small entity, if applicable. A Small Entity Statement must also be filed (Note 37 C.F.R. 1.9, 1.27, 1.28).</p>					
SUBTOTAL = <input type="text" value="\$ 1598.00"/>					
<p>Processing fee of \$130.00, for furnishing the English Translation later than <input type="checkbox"/> 20 <input type="checkbox"/> 30 months from the earliest claimed priority date (37 C.F.R. 1.492(f)).</p>					
+ <input type="text" value="0.00"/>					
TOTAL NATIONAL FEE = <input type="text" value="\$ 1598.00"/>					
<p>Fee for recording the enclosed assignment (37 C.F.R. 1.21(h)). The assignment must be accompanied by an appropriate cover sheet (37 C.F.R. 3.28, 3.31). \$40.00 per property + <input type="text" value="\$ 0.00"/></p>					
<p>Fee for Petition to Revive Unintentionally Abandoned Application (\$1240.00 - Small Entity = \$620.00) <input type="text" value="\$ 0.00"/></p>					
TOTAL FEES ENCLOSED = <input type="text" value="\$ 1598.00"/>					
<table border="1" style="width: 100%;"> <tr> <td style="width: 60%;">Amount to be:</td> <td style="width: 20%; text-align: center;">Refunded <input type="checkbox"/></td> <td style="width: 20%; text-align: center;">Charged <input type="checkbox"/></td> </tr> </table>			Amount to be:	Refunded <input type="checkbox"/>	Charged <input type="checkbox"/>
Amount to be:	Refunded <input type="checkbox"/>	Charged <input type="checkbox"/>			
<p>a. <input checked="" type="checkbox"/> A check in the amount of \$ 1598.00 to cover the above fees is enclosed.</p> <p>b. <input type="checkbox"/> Please charge my Deposit Account No. 14-1140 in the amount of \$ <input type="text"/> to cover the above fees. A duplicate copy of this form is enclosed.</p> <p>c. <input checked="" type="checkbox"/> The Commissioner is hereby authorized to charge any additional fees which may be required, or credit any overpayment to Deposit Account No. 14-1140. A <u>duplicate</u> copy of this form is enclosed.</p> <p>d. <input type="checkbox"/> The entire content of the foreign application(s), referred to in this application is/are hereby incorporated by reference in this application.</p>					
<p>NOTE: Where an appropriate time limit under 37 C.F.R. 1.494 or 1.495 has not been met, a petition to revive (37 C.F.R. 1.137(a) or (b)) must be filed and granted to restore the application to pending status.</p>					
SEND ALL CORRESPONDENCE TO:		<u>H. Warren Burnam, Jr.</u> SIGNATURE			
<p>NIXON & VANDERHYE P.C. 1100 North Glebe Road, 8th Floor Arlington, Virginia 22201 Telephone: (703) 816-4000</p>					
H. Warren Burnam, Jr. NAME					
29,366 October 19, 2000 REGISTRATION NUMBER Date					

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of

RUMSEWICZ et al

Atty. Ref.: 100-43

Serial No. to be assigned

Group: unknown

Filed: October 19, 2000

Examiner: unknown

For: Access Control Method and Apparatus

* * * * *

Assistant Commissioner for Patents
Washington, DC 20231

Sir:

PRELIMINARY AMENDMENT

In order to place the above-identified application in better condition for examination, please amend the application as follows:

IN THE CLAIMS

Claim 3, line 1, delete "or claim 2".

Claim 4, line 1, change "any one of claims 1 to 3" to -- claim 1 -- .

Claim 7, line 1, delete "or claim 6".

Claim 10, line 1, delete "or claim 9".

Claim 13, line 1, change "any one of claims 4 to 12" to -- claim 4 -- .

Claim 14, line 1, change "any one of claims 4 to 12" to -- claim 4 -- .

Claim 18, line 1, delete "or claim 17".

Claim 20, line 1, change "any one of claims 16 to 19" to -- claim 16 -- .

Claim 24, line 1, change "any one of claims 16 to 23" to -- claim 16 -- .

Claim 27, line 1, delete "or claim 26".

Claim 32, line 1, change "any one of claims 28 to 31" to -- claim 28 -- .

Claim 34, line 1, change "any one of claims 16 to 33" to -- claim 16 -- .

Claim 35, line 1, change "any one of claims 24 to 34" to -- claim 24 -- .

Claim 36, line 1, change "any one of claims 24 to 34" to -- claim 24 -- .

Claim 44, line 1, change "any one of claims 24 to 43" to -- claim 24 -- .

Claim 45, line 1, delete "any one";

line 2, change "of claims 16 to 44" to -- claim 16 -- .

Claim 46, line 1, delete "any one";

line 2, change "of claims 16 to 44" to -- claim 16 -- .

REMARKS

The above amendments are made to place the claims in a more traditional format.

Respectfully submitted,

NIXON & VANDERHYE P.C.

October 19, 2000

By: H. Warren Burnam

H. Warren Burnam, Jr.
Reg. No. 29,366

HWB:lsh
1100 North Glebe Road, 8th Floor
Arlington, VA 22201-4714
Telephone: (703) 816-4000
Facsimile: (703) 816-4100

ACCESS CONTROL METHOD AND APPARATUS

The present invention relates to a method and apparatus for controlling access to one or more servers in a network, and more particularly relates to a method and 5 apparatus that controls the admission of new users to the one or more servers such that users already accessing server resources will receive priority service over new users attempting to access a server.

In a computer network, such as the Internet which represents a vast number of computers linked to each other worldwide, information exchange such as e-mail may 10 occur between various computers linked through the Internet using a set of protocols such as the transmission control protocol/Internet protocol (TCP/IP). A "client" may typically be a web browser running on a computer. A client typically communicates with a "server". A server may uniquely distinguish clients using an IP address and possibly other identifiers. Other examples of clients include proxy servers and web 15 robots. Information on the Internet may be made available to the public through "servers". Servers are computers which make available files or documents they hold. The worldwide web (www) is a method of accessing this information from the servers and allows a user to navigate the resources of the Internet by displaying or downloading pages of information that are stored at the servers. Clients communicate 20 with servers using the hypertext transfer protocol (HTTP). In the Internet the HTTP protocol is generally transmitted over TCP/IP.

Usage of the Internet and the worldwide web is increasing at a very rapid rate and more and more new users are being connected to the Internet and www. As its 25 growth continues, the www will provide a rapidly growing number of commercial services, with applications ranging from information retrieval for text, images or multi media through to purchasing, for example, ticket and item sales. As such, there is a need to improve the reliability of the web so as to make it a suitable medium for high volume, business critical applications.

During periods of high use of a web server, even in situations such as 30 overload of the server wherein the load of requests to the web server exceeds the capabilities of the server, the server still allows requests from new users. These capabilities include the physical capacity of the CPU, the memory and the network.

Continuing to allow requests from new users tends to degrade the performance perceived by users that are already accessing the server system and which is evidenced through long delays or inability to service information requests. For example users already accessing web pages on a server may experience poor 5 performance as they navigate through links to pages on that same server. The performance as perceived by already accessing users is assessed in terms of maintaining throughput, response time levels and system stability rather than by increasing the overall throughput of the web server.

The present invention seeks to address this problem by providing a system 10 that allows a server to accept or deny access by new users to that server in preference to those already accessing the server. Such a decision for accepting or rejecting the admission of new users is based on certain performance parameters, such as the setting of the congestion level at the server.

Accordingly, the present invention provides a method of controlling access to 15 a server by a client to a server in a network, said method comprising the steps of:

monitoring resource usage of said server;
allowing a connection of said client to pass to said server where said connection forms part of an active session wherein said connection forms part of said active session if at least one previous connection from said client has been passed to said 20 server within a predetermined time interval, and

allowing or rejecting a new connection of another client to pass to said server according to an admission control scheme.

The step of allowing or rejecting may be based on resource usage of said server. The method may include searching for an active session of said client and 25 admitting the client's connection to pass to said server if there is an active session for said client.

The method may include determining a congestion level from monitoring the resource use on said server. The congestion level may be represented by a whole 30 number between zero and c_{max} where zero indicates that the server is operating normally and c_{max} is the maximum congestion level. The congestion level may be used to calculate an admission interval which represents the time after which a new session for a new connection can begin.

A new connection may be admitted after the admission interval expires. Periodically, a determination is made on whether the resource usage is high, and if it is high the congestion level is increased. The method may include checking whether a client has had a connection attempt rejected within a predetermined period.

5 The present invention also provides an interface unit for controlling access to a server by a client in a network, said interface unit comprising:

means for monitoring resource usage of said server;

admission means for allowing a connection of said client to pass to said server where said connection forms part of an active session wherein said connection forms 10 part of said active session if at least one previous connection from said client has been passed to said server within a predetermined time interval,

said admission means further rejecting or allowing a new connection of another client to pass to said server according to an admission control scheme.

The admission means further rejecting or allowing said new connection may 15 be based on said resource usage of the server. The monitoring means may be in the form of a storage means, such as a database that maintains information about previous sessions, a list of active sessions and routinely polls said server to obtain resource information on that server. The interface unit may be linked to the computer network through an external port and be linked to the server through an internal port. 20 The interface unit may be a computing processor and the admission means may be a control unit for rejecting or admitting a new connection to said server.

The invention will now be described in a preferred embodiment with reference to the accompanying drawings wherein:

Figure 1 shows a computer network system that links clients to servers 25 through the network:

Figure 2 is the system of Figure 1 showing just one interface unit linking one server to a number of clients through the network:

Figure 3 is a diagram showing the individual components of a server and interface means.

30 Figure 4 is a flow diagram of a first connection admission control scheme;

Figure 5 is a flow diagram of a second connection admission control scheme;

Figure 6 is a detailed view of one decision point in the flow diagram of Figure 4;

Figure 7 is a detailed view of one decision point in the flow diagram of Figure 5.

5 Figure 8 shows the flow of message streams through the interface means between servers and clients;

Figure 9 is a flow diagram of a first implementation of the interface means;

Figure 10 is a diagram of buffer and stream management for the implementation of Figure 9, and

10 Figure 11 is a process diagram of a second implementation of the interface means.

In Figure 1 there is shown a system 2 that links one or more clients 4, either being workstations or PCs, with one or more servers 6 through a network 8, which may be a computer network such as the Internet. Each of the servers 6 is connected 15 to the computer network 8 through interface means 10. Typically when a user wishes to access a server 6 for the purpose of reading or downloading a web page, they will initially input the address or universal resource locator (URL) identifying a web site or web page that one or more server 6 will handle or have stored. Where one of the servers 6 has many requests to obtain information for a number of users it can 20 become overloaded and consequently provide a degradation in service to each one of the users accessing the server 6 through respective clients 4. In situations where the demands on a server 6 are high, even on the verge of overloading the server 6, the interface means 10 which is implemented as a computing processor, and more particularly as an interface unit, monitors critical resource usage on the server based 25 on predefined parameters, and maintains a derived parameter known as the congestion level of the system. Based on this congestion level, the interface means 10 controls a number of new users allowed to access data on the server 6. At this stage each of these new users are defined as "not participating in an active session" on server 6. Thereafter any one of these new users that is allowed to access data is 30 then said to have an "active session" on the server 6. All other users are blocked and are said to "not have an active session" on the server 6. The blocked users are either placed in a queue until such time they are allowed access to the data or the users are

rejected outright. Users placed in the queue may also be eventually rejected outright. The definition of "active session" generally is that the user has communicated with a server within a defined period, typically the last N minutes. This definition is based on user behaviour, for example, where a user downloads a page, spends a short 5 period of time reading the page, then downloads another page and so on. The interface unit 10 also determines, in real-time, the rate at which new requests to the server 6 can be accepted while still providing an acceptable quality of service for already accepted requests. It also maintains a table of "active sessions", and accepts an appropriate rate of new requests while rejecting all other new requests.

10 The interface unit 10 implements an admission control scheme to determine whether users are admitted or rejected. This determination is based on a current and past server resource use at the web server 6 such as memory capacity, bandwidth use on the link from the server 6 to the Internet 8, resources available on the interface unit 10 and processor load or utilisation. Specifically, server resources that are 15 monitored include CPU utilisation, number of pages, number of pages swapped in and out, memory used, number of disk input/output operations and average disk delay since the last poll by the interface unit 10 of the server 6 for its current resource usage took place. If sufficient resources are not available, a rejection message (possibly with an explanatory note) is sent to users who attempt to gain admission but 20 are not admitted.

As mentioned previously, any one or more of the servers 6 communicates with each of the clients 4 through the interface unit 10 over the computer network 8. The interface unit 10 may be connected to several servers 6 and determines admission of new sessions by monitoring resource usage on each server 6. A session 25 with a client represents at least one admitted connection from the client. Referring to Figure 2, the interface unit 10 is linked to the Internet 8 through link 12 and to the server 6 through link 14. Generally it is preferred that the bandwidth of the link 14 between interface unit 10 and server 6 exceeds the bandwidth of link 12 between the interface unit 10 and the computer network or the Internet 8.

30 The architecture Figure 2 can be realised in several different ways. Firstly, the interface unit 10 can run on a multihomed computer with one connection running to the Internet 8 and another connection to the computer that supports the server 6. A

multihomed computer has two or more physical interfaces to Internet protocol based networks. Secondly, the interface unit 10 can run on the computer supporting the web server 6 with the interface unit 10 using Internet domain sockets to communicate with the Internet 8 and Unix domain sockets to connect to the web server 6. The 5 HTTP protocol is used on both the Unix domain and the Internet domain sockets. Thirdly, the interface unit can be built into the web server 6 so that all communication with the Internet 8 is performed by the interface unit module and communication with the server component is implemented through either function calls or shared memory.

10 In Figure 3 there is shown the components of each of the interface unit 10 and the web server 6. An external port 16 of the interface unit 10 is a port supporting the HTTP protocol that connects the interface unit 10 to the Internet 8. An internal port 18 of the interface unit 10 is a port connecting the interface unit 10 to the www server 6. The ports are bi-directional. The number of outstanding connection requests 15 which a port can maintain may be limited. Connections are stored in a buffer associated with external port 16 which must be polled to extract data. After polling, the IP address of the connection's client is then known. Data must be placed in a buffer associated with the external port 16 to be sent. The IP address is a unique identifier of a computer connected to the Internet 8 and each client is identified by 20 such a unique IP address possibly in connection with further information.

 A connection between the interface unit 10 and the client 4 is said to exist from when the interface unit 10 polls the port for the connection until the connection is closed. Data can be transferred bi-directionally through a connection. A client may begin an unrestricted number of connections, for example, a page request from a 25 web browser typically requires a connection to receive the text and a connection for each graphical image. Under HTTP 1.0 a connection is closed by the server 6 immediately after the information has been transferred. Connections may be closed immediately by the server 6 and a connection does not imply admission to the server 6. Connected between the external port 16 and internal port 18 is a control unit 20 30 that admits or rejects a connection to the server 6. Control unit 20 admits or rejects new users in accordance with an admission control scheme to be described. A storage means 22, in the form of a history data base, implemented as a memory unit,

maintains current and recent information for determining admission of connections such as:

- resource usage,
- information about previous sessions and their resource usage so that predictions may be made about a future session,
- a list of active sessions of clients who currently have access to one or more servers, and
- in the case of a percentage throttling admission scheme (to be described), a history of rejected clients is also stored.

10 An active session is said to exist if the last connection from a client was made within a specific time interval, which is likely to represent the time period necessary to read a web page. A session ends when no connections are received from the client for the specified time interval. An active session must be associated with a history entry and a client.

15 When a new client connection is accepted by the interface unit 10 for the server 6, an entry is added to an "Active Session Table" in the history database 22. The entry includes the time at which the last client-server interaction occurred. This interaction can be either a user request arriving or the transmission of information from the server to the client. If more than INACTIVE_SESSION_TIMEOUT 20 seconds elapses between server-user interactions, the session is deemed inactive. This is an indication that the user is no longer interested in the server 6. Inactive sessions are removed from the Active Session Table.

25 The history data base 22 obtains information on resources of the server 6 from a server storage means 24, in the form of a resources data base, over link 26. The server storage means 24 resides within the server 6. Every STATUS_REQUEST_TIME seconds, the database 22 polls the resources database 24 for that server's current resource usage. A server processor 28 handles all of the requests from various users for obtaining web pages and the like and is linked to the internal port 18 of the interface unit 10 through link 14.

30 The operation of the interface unit 10 will now be described with reference to an admission control scheme called call gapping admission. With reference to Figure 4 there is shown a flow diagram of the processes involved using the call gapping

admission control scheme. The processes of the server 6 operate in a continual loop and are divided up into two sections, one being for the admission control part 402 the other being for the data transfer and history updating part 404. At step 406 the interface unit 10 listens or checks for connections from any one of the clients 4 at the 5 external port 16. The external port 16 may store connection requests in a queue buffer or a connection request queue for a limited time. The external port 16 must be checked often for connections to avoid loss of connections including those for active sessions. If there are no connections from clients, flow proceeds to step 424 of the data transfer and history updating part 404. If there is a connection the process 10 moves to step 408 where the history data base 22 is queried for an active session having the same IP address as the client making the connection. Only one connection can be handled for each iteration of the outside loop. If there is an active session, the connection is accepted by the control unit 20 and provided that there is an available "connection slot" at step 416. A connection slot is required for the 15 transfer of data through a connection such that the slot maintains information relevant to the connections such as buffers, buffer sizes, the IP address associated with the external port buffer and the file handles for the connection. Data is stored in the internal port buffer prior to sending. If there is no available connection slot the connection is closed at step 422. If no data (including a rejection message) has been 20 sent from the external port 16 to the client 4, the connection has been ignored. If there is a connection slot available the process proceeds to the data transfer and history updating part 404 starting at step 424. If the connection is a new connection it will only be added at step 418 if it meets the criteria of the call gapping admission scheme. A new connection is a connection from a client for which the history 25 database contains no information. A new connection originates from a client who is "currently not participating in an active session" with server 6. The admission of a new connection begins an active session. The new connection is then placed in a queue at step 419 until such time that the data is able to be transferred. The call gapping admission scheme allows one new session to be started only after a 30 calculated time interval has elapsed since the last new session began. This time interval is determined by the current congestion level. The congestion level is a measure of how much traffic or requests are initiated from the clients 4 and is

determined by monitoring resource use on one or more servers 6 and the utilisation of the link 12 to the Internet 8. A congestion level is represented by a whole number between zero and c_{max} inclusive where c_{max} is the maximum congestion level, with zero meaning that the server is operating normally. Therefore a high resource use 5 increases the congestion level while a low resource use reduces the congestion level. Thus, when a request or connection arrives, a search is made of the Active Session Table. If the connection comes from an active session, it is immediately passed through to server 6. This gives priority to active sessions over new sessions.

Going back to step 408, if no active session is found through the search then 10 the process proceeds to step 410 where a search is conducted for a "history entry" or an expired one that may be recycled if it exists. A history entry stores information about a client such as the last time that a connection from a client was closed and the IP address of the client. History entries may be recycled if they do not represent an active session. If there is no empty history entry the process moves to step 422 where 15 the connection is closed. If an empty history entry is found or an expired one is recycled, then the process moves to step 412 and the connection is either admitted or rejected according to the call gapping admission scheme. For a connection to be admitted or for a rejection message to be sent, there must be a history entry available. If not, the connection is closed immediately.

20 With reference to Figure 6 the processes involved in admitting or rejecting a connection is disclosed in greater detail. Firstly at step 413 the admission interval is calculated. This is the interval after which a new session may begin and is a function of the congestion level. Firstly $f(c)$ is calculated where c is the congestion level and f is a function which determines the time interval given the congestion level c . The 25 process then moves to step 415 where an admission or rejection determination is made. A new connection will be admitted if

$$t - t_1 > f(c)$$

where t is the current time, t_1 is the time of the last admission. If this is true then the connection is admitted and the process moves to step 417 where the last admission 30 time is set. This is where the time of the last connection admission is updated to the current time, in other words,

$$t_1 = t.$$

If the connection is admitted, a new active session begins provided a connection slot is available at step 416. A new history entry is then entered in the history data base 22. If the new connection is rejected the process moves to step 414 where a check is made for a connection slot, and if one is available, then a rejection message is added 5 to the external port buffer for a specific client at step 420. If no connection slot is available then the connection is closed at step 422. The process reaches step 424 where a check is made to see if a constant time interval I has elapsed such that

$$t - t_u > I$$

where t is the current time and t_u is the time of the last update. The process then 10 moves on to step 426 to determine if the resource use of the server 6 is high or if any of the resources are above predefined threshold levels. If this is the case, then the congestion level c_n is increased at step 430. The congestion level c_n is defined as

$$c_n = \min (c_{n-1} + 1, c_{\max})$$

where c_n is the current congestion level, c_{n-1} is the previous congestion level and c_{\max} 15 is the maximum congestion level. So the congestion level is essentially incremented by one value. If the resource usage is determined not to be above predefined levels (not high), the process moves to step 428 to determine if the resource usage is low or in other words if all resources are below predefined threshold levels. If this is the case then the process moves to step 432 where the congestion level is decreased. The 20 congestion level c_n in this case is then defined as

$$c_n = \max (c_{n-1} - 1, 0)$$

where c_{n-1} is the previous congestion level. Therefore the congestion level c_n is 25 decremented by one value. If both the resource usage is neither high nor low or somewhere in between the predefined levels, then the process moves to 434 where updates take place in the history data base 22. This also occurs where an increase or decrease in the congestion level has occurred from steps 430 and 432 respectively.

The updates that occur include other history information, besides congestion level, for resources such as the capacity of the link 14 and the interface unit 10. This information is stored as entries in a database which correspond to regular, specified 30 time intervals. The values held by the control unit 20 are stored in the database 22 if the time interval I has elapsed since the last update. The time is then set to $t = t_u$ where t is the current time and t_u is the time of the last update.

Once the update has taken place at step 434 the process moves to step 436 where data is copied from the internal port 18 to the external port 16 of the interface unit 10. For each connection, if an error occurs or there is no remaining data to be sent the connection is closed. Where the interval I has not elapsed at step 424 if there 5 is any data to be transferred between the ports 16 and 18 this takes place at step 436 as well.

All of the rejection messages that have been stored at step 420 in a message buffer, associated with the external port, are sent to the respective clients at step 438 and if there is no information remaining in the message buffer, then the connection is 10 closed for each client. The process then reverts back to step 406.

Essentially the purpose of the data transfer and history updating unit 404 is to maintain the history information about the link 14 between the interface unit 10 and the server 6 and the resource usage of the server or servers 6. It is also used to transfer data between both the internal and external ports 18, 16 of the interface unit 15 10 and to send the rejection messages via the external port 16. The history information and congestion level are only updated if a specified time interval I has elapsed since the last update.

Another admission control scheme is known as the percentage throttling admission scheme which allows a proportion of new sessions to be started according 20 to the particular congestion level. A flow diagram depicting this scheme is shown at Figure 5 which is split into two parts, one part being for admission control 502 and the other part being for data transfer and history updating 504. All of the steps shown in the flow diagram in Figure 5 are identical to those shown at Figure 4 except for a new step 509 and an altered admit step 512, both of which will be described 25 hereinafter. Otherwise steps 506 through to 538 correspond to steps 406 to 438. As a result of the search query at step 508, if no active session, given the IP address of a connection, is found then the process moves to step 509 where a search is undertaken for any rejected history or session. This is a check to see whether a specific client has been rejected within a specified time period. If a recent session or recent history has 30 been rejected then the connection is closed immediately at step 522, otherwise the process moves to step 510 where a search is made for an empty history entry or one that may be recycled.

This additional decision action at step 509 addresses the problem of clients gaining admission by frequently making connections. Hence, whenever a connection arrives, it passes through a filter at this step 509. If the connection arrives from a client that has recently had a connection rejected, then the connection is closed 5 without sending a rejection message.

The admit step at 512 is shown in more detail in Figure 7. After step 510 the process moves to step 512 where at step 513 a calculation is made for the rejection proportion. This is the proportion of new connections to be rejected and is

$$p_r = \frac{c}{c_{\max}}$$

10 where p_r is the proportion of connections to be rejected, c is the congestion level and c_{\max} is the maximum congestion level. At step 515 an admission determination is made to admit or reject a new connection. This will happen where $p_r < r$ where r is a random number between 0 and 1 inclusive, which is generated with each new connection. If the equation holds true the new connection is admitted and the process 15 moves to step 516, otherwise the new connection is rejected and the process moves to step 514.

The implementation of the interface unit 10 has been performed in two ways. Firstly, a C implementation that maintains the active session table based only on the IP - address of the source of the request, and secondly using an Erlang based system 20 which uses an active session table based on "cookies" to provide individual user identification for admission control. Cookies are a small quantity of data sent from a web server to a client and stored by the client. The value of a cookie can be requested by a web server.

The interface unit 10 has been developed on a Sun Ultra Enterprise 2 25 SparcServer, running Solaris 2.5.1 and Intel Pentium based computers running FreeBSD 2.2.2. It is assumed that the web server 6 has the Unix based "rstat" facility.

Shown in Figure 8 there is a revised version of Figure 2 showing the message 30 streams between the various components. The system 2 is separated into two networks, an inside network, comprising the servers 6 and a client side of the

interface unit 10, and an outside network comprising a server side of the interface unit 10 together with the Internet 8 and web clients 4. To the outside network, the interface unit 10 appears as a web server. Clients 4 located on the Internet 8 may query the interface unit 10 to gain access to pages held by any one of the web servers

5 6. When a user wishes to view a page held on the web server 6 their web client 4 opens a stream S1 and issues request A to the interface unit's server side. Admission control is performed on this request using the IP address of the client or a cookie as the key. If the request is accepted, a stream S2 is opened to the web server 6 and the request is forwarded to the web server on link C. Otherwise a rejection message is

10 sent to the client through link B via the stream S1 to the client. The web server 6 responds to the request C using the stream S2 to the interface unit's client side and responds with the requested page D. The requested page is forwarded to the web client 4 through E via the open stream S1 over the Internet 8.

In normal operation all data read from web client 4 is copied to the web server

15 6 and all data from the web server 6 is copied to the web client 4. However, a delay can be introduced to allow the contents of a data stream to be monitored by buffering the flow of data between the interface unit's server side and the web client 4 to determine the nature of the request, and between the interface unit's client side and the web server 6 to recover HTTP status codes. Furthermore, data can be injected

20 into the flow to allow a replacement page to be sent to the web client 4 or to augment an existing page.

A number of statistics may be gathered by the interface unit 10 including the number of connection attempts, web server errors, the number of bytes transferred and the transfer rate. Additional statistics are gathered from the web server 6 using

25 the "rstat" facility. There are two major versions of the HTTP protocol existing. HTTP/1.0 is the most common and opens and closes the streams for each file transferred from the web server 6. HTTP/1.1 is a newer protocol that opens a stream for a session and allows multiple requests and responses to be transferred through the streams. The current versions of the interface unit 10 are designed to handle both

30 protocols. The C version implementation of the interface unit 10 will now be described.

The C implementation is implemented as a single process which continuously loops through the actions of checking for new connections, admitting or rejecting offered connections and then servicing current connections by transferring bytes to and from the two buffers associated with the external and internal ports 16 and 18.

5 All memory including tables, history data and buffers is either static or allocated at start up. No dynamic memory operations occur in the main loop of the server. This simplifies the design and eliminates the need to handle the non-availability of a dynamic resource.

With reference to Figure 9, the process 600 starts at step 602 where a listen 10 socket is opened to check for new connections and existing connections. At step 604 the connections are accepted and then the process moves to step 606 where a decision is made as to whether the connection is pending. If it is, it goes to step 608 where another decision is made to admit that connection otherwise if the connection isn't pending it goes to step 616 (to be described). At step 608, if the connection is not 15 admitted it goes to a queue rejection at step 610 which in turn then goes to step 616. Otherwise, if the connection is admitted then the connection is opened to the web server 6 at step 612. Once the connection is opened a record is inserted into the active connection table at step 614. When the connection is accepted by the interface unit 10 a stream is opened up to the web server 6. The file handles of the stream S1 20 from the client 4, as seen in Figure 10, and the stream S2 to the web server, are recorded in an empty slot in the active connection table. The active connection table also contains two buffers for bytes travelling from the web server and from the web client respectively as seen in Figure 10 also. These buffers are used to store incoming bytes read from a stream. If there is data in a buffer it is copied to the 25 output stream and if there is a space in the buffer then data is read from the input stream. At step 616, periodically the process scans through the active connection table and the following copying processes take place:

1. For each non-empty entry data is copied from the operating system buffers on the server side of the interface unit 10 into the "buffer of bytes from the web client" 4;
2. Data is copied from the "buffer of bytes from the web client" 4 to the web server" 6;

3. Data is copied from the operating system buffers on the interface unit's client side into the "buffer of bytes from the web server" 6; and
4. Data is copied from the "buffer of bytes from the web server" 6 to the web client 4.

5 The second implementation, the Erlang version, was developed and run using Erlang OTP R3A. The Erlang version consists of a collection of co-operating Erlang processes as to be described with reference to Figure 11. The initial function call creates a listening socket at step 702 and a server process is spawned at step 704 which is repeated on the right hand side. The server process then waits on an accept 10 until a web client 4 opens a stream S1 to the interface unit's server side. The stream is accepted at step 706 and another copy of the server is immediately spawned at step 708. It is essential that a copy of the server process is always waiting in the accept state to ensure that the interface unit 10 is always able to handle a new request. Admission control is then performed at 710 and if the connection is not to be 15 admitted then a rejection message is sent back to the client at step 712. If the connection is to be admitted, the Xfer routine is called at step 714. The Xfer routine is shown on the right hand side of Figure 11. The connection is then opened to the web server at step 716 and data is copied from the client to the server and from the server to the client, or in other words between the web server and the web client at 20 step 718 and then at step 720 when both client and server sockets are closed, the process is stopped at step 722. If both sides of the connection are still open then the data copying process is still occurring at step 718.

25 The Erlang implementation is more elegant than the C version. Using separate processes to handle each active connection allows code to be simplified and eliminates the need for a table of active connections. Once started, the Erlang processes perform the necessary data transfers and report their statistics without requiring oversight.

30 Statistics are gathered from the web server 6 using the rstat mechanism and a piece of C glue code. The glue code makes a copy of the structure returned by rstat and converts it into an Erlang tuple. An Erlang process periodically polls the rstat service via the glue code.

The system and method of the present invention allows a user, where they have access to a web server, to receive priority access to that server's resources over new users. This ensures that delay in moving from web page to web page on that server is minimised thereby ensuring that the service the user perceives is maintained 5 at an acceptable level.

Throughput is maintained at the highest possible level by ensuring that the interface unit only controls traffic to the web server when the servers resources have, or are approaching, exhaustion. Therefore by controlling traffic before resources are exhausted on a server, the chances of the server breaking down are minimised.

10 As an example, in a commercial environment such an idea is important when many customers of a company attempt to access the company's web home page to order products. Typically they will go to the company home page and click on a "products" link, and then possibly an "order form" link. If each of these moves is very slow, the customer is likely to be dissatisfied with the level of service received., 15 particularly in the light of the time wasted.

By employing the system of the present invention, the company in the above example ensures that accepted users receive good levels of service while other users are requested to return later when the server is less busy, thereby minimising the amount of time they waste.

CLAIMS:

1. A method of controlling access to a server by a client in a network, said method comprising the steps of:
 - 5 monitoring resource usage of said server;
 - allowing a connection of said client to pass to said server where said connection forms part of an active session wherein said connection forms part of an active session if at least one previous connection from said client has passed to said server within a predetermined time interval; and
 - 10 allowing or rejecting a new connection of another client to pass to said server according to an admission control scheme.
2. A method according to claim 1 wherein the step of allowing or rejecting is based on said resource usage of said server.
- 15 3. A method according to claim 1 or claim 2 further comprising the step of searching for an active session of said client and admitting a connection of said client to pass to said server if there is an active session for said client.
- 20 4. A method according to any one of claims 1 to 3 wherein said admission control scheme includes the step of determining a congestion level in response to said monitoring step.
5. A method according to claim 4 further comprising the step of calculating a further predetermined time interval based on said congestion level, whereby said new connection is allowed to pass to said server after said further predetermined time interval has expired.

6. A method according to claim 5 wherein said further predetermined time
interval represents the time since a new connection was allowed.
7. A method according to claim 5 or claim 6 wherein the allowance of said new
connection begins an active session.
5
8. A method according to claim 4 further including the step of allowing a
proportion of new connections to pass to said server based on said congestion
level.
10
9. A method according to claim 8 further comprising the step of searching for
previous rejections or previous connections whereby a new connection of a
client is rejected and closed if said client has had a previous connection or a
previously rejected connection within a specified time period.
15
10. A method according to claim 8 or claim 9 wherein said proportion of new
connections to be rejected p_r is
20

$$p_r = \frac{c}{c_{\max}}$$

where c is the congestion level and c_{\max} is the maximum congestion level.

11. A method according to claim 10 further including the step of generating a
random number r where r is between zero and one inclusive for each new
connection.
25
12. A method according to claim 11 wherein new connections are allowed to pass
to said server if $p_r \leq r$ or are rejected if $r < p_r$.
30

13. A method according to any one of claims 4 to 12 further including the step of increasing the congestion level if the resource usage is above a first predetermined threshold level.
- 5 14. A method according to any one of claims 4 to 12 further including the step of decreasing the congestion level if the resource usage is below a second predetermined threshold level.
- 10 15. A method according to claim 4 wherein said congestion level is represented by a number between zero and c_{\max} , where zero indicates that the server is operating normally and c_{\max} is the maximum congestion level.
16. An interface unit for controlling access to a server by a client in a network, said interface unit comprising:
 - 15 means for monitoring resource usage of said server;
 - admission means for allowing a connection of said client to pass to said server where said connection forms part of an active session wherein said connection forms part of an active session if at least one previous connection from said client has passed to said server within a predetermined time interval;
 - 20 said admission means further rejecting or allowing a new connection of another client to pass to said server according to an admission control scheme .
- 25 17. An interface unit according to claim 16 wherein said admission means further rejecting or allowing said new connection is based on said resource usage of said server.
18. An interface unit according to claim 16 or claim 17 wherein said interface unit is connected to said computer network through a first port and connected to said server through a second port.

19. An interface unit according to claim 18 wherein said admission means is connected between said first port and said second port.
- 5 20. An interface unit according to any one of claims 16 to 19 wherein said monitoring means is a storage means, such as a database, said storage means linked to said admission means.
- 10 21. An interface means according to claim 20 wherein said storage means maintains a table of active sessions of respective clients.
22. An interface unit according to claim 21 wherein said table of active sessions is scanned to determine if said connection is associated with an active session in said table of active sessions.
- 15 23. An interface unit according to claim 22 wherein if an active session exists in said table of active sessions for said connection, said connection is passed to said server through said admission means.
- 20 24. An interface unit according to any one of claims 16 to 23 wherein said admission control scheme includes setting a congestion level.
- 25 25. An interface unit according to claim 24 wherein a further predetermined time interval is calculated based on the congestion level such that a new connection is allowed to pass to said server after said further predetermined time interval has expired.
26. An interface unit according to claim 25 wherein said further predetermined time interval represents the time elapsed since a new connection was allowed.

27. An interface means according to claim 25 or claim 26 wherein said new connection is assigned an active session which is subsequently stored in said table of active sessions.

5

28. An interface unit according to claim 24 wherein said admission means allows a proportion of new connections to pass to said server based on said congestion level.

10 29. An interface unit according to claim 28 wherein said proportion of new connections to be rejected p_r is

$$p_r = \frac{c}{c_{\max}}$$

15 where c is the congestion level and c_{\max} is the maximum congestion level.


20 30. An interface unit according to claim 29 wherein a random number r is generated for each new connection where r is a number between zero and one inclusive.



31. An interface unit according to claim 30 wherein said admission mean allows new connections to pass to said server if $p_r \leq r$ or are rejected if $r < p_r$.


25 32. An interface means according to any one of claims 28 to 31 wherein said storage means stores information about previous connections including previously rejected connections of various clients.


33. An interface unit according to claim 32 wherein a new connection of a client is rejected and closed if the client has had a previous connection or a previously rejected connection within a specified time period.
- 5 34. An interface unit according to any one of claims 16 to 33 wherein said monitoring means periodically polls said server to obtain information on said resource usage of said server and subsequently stores said resource usage information.
- 10 35. An interface unit according to any one of claims 24 to 34 wherein said congestion level is increased if said resource usage is above a first predetermined threshold level.
- 15 36. An interface unit according to any one of claims 24 to 34 wherein said congestion level is decreased if said resource usage is below a second predetermined threshold level.
37. An interface unit according to claim 23 wherein said connection is passed to said server using a connection slot.
- 20 38. An interface unit according to claim 22 wherein if no active session exists in said table of active sessions, a search is performed of said storage means for an empty history entry.
- 25 39. An interface unit according to claim 38 wherein if said empty history entry is found or an expired history entry is found, said connection is either allowed or rejected in accordance with said admission control scheme.

40. An interface unit according to claim 38 wherein if said history entry is not found, said connection is subsequently closed.
- 5 41. An interface unit according to claim 39 wherein an allowed connection is passed to said server using a connection slot or a rejected connection is returned to the respective client using a connection slot and forwarded with a rejection message.
- 10 42. An interface unit according to claim 18 wherein said first port has buffer means for storing connections prior to being allowed or rejected by said admission means.
- 15 43. An interface unit according to claim 42 wherein said first port is routinely polled for stored connections.
44. An interface unit according to any one of claims 24 to 43 wherein said storage means is updated after a time interval I with information including resource usage of said server, history information and the present congestion level.
- 20 45. A communications system incorporating an interface unit according to any one of claims 16 to 44 wherein said interface unit is connected between said server and said network and one or more clients are linked to said server through said network.
- 25 46. A communications system incorporating an interface unit according to any one of claims 16 to 44 wherein said interface unit forms part of said server and said interface unit is linked to one or more clients through said network.














III-4B





5A












III 10.



RULE 63 (37 C.F.R. 1.63)

INVENTORS DECLARATION FOR PATENT APPLICATION
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

As a below named inventor, I hereby declare that my residence, post office address and citizenship are as stated below next to my name, and I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

ACCESS CONTROL METHOD AND APPARATUS

the specification of which (check applicable box(s)):

is attached hereto
 was filed on October 19, 2000 as U.S. Application Serial No. PCT/AU99/00297 (Atty Dkt. No. 100-43)
 was filed as PCT International application No. on 20 April 1999
and (if applicable to U.S. or PCT application) was amended on

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above. I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with 37 C.F.R. 1.56. I hereby claim foreign priority benefits under 35 U.S.C. 119/365 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed or, if no priority is claimed, before the filing date of this application:

Priority Foreign Application(s):

Application Number	Country	Day/Month/Year Filed
PP 3082	Australia	20 April 1998

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below.

Application Number	Date/Month/Year Filed
--------------------	-----------------------

I hereby claim the benefit under 35 U.S.C. 120/365 of all prior United States and PCT international applications listed above or below and, insofar as the subject matter of each of the claims of this application is not disclosed in such prior applications in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. 1.56 which occurred between the filing date of the prior applications and the national or PCT international filing date of this application:

Prior U.S./PCT Application(s):

Application Serial No.	Day/Month/Year Filed	Status: patented pending, abandoned
PCT/AU99/00297	20 April 1999	

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon. And on behalf of the owner(s) hereof, I hereby appoint NIXON & VANDERHYE P.C., 1100 North Glebe Rd., 8th Floor, Arlington, VA 22201-4714, telephone number (703) 816-4000 (to whom all communications are to be directed), and the following attorney(s) thereof (or of the same address) individually and collectively owner(s') attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith and with the resulting patent: Arthur R. Crawford, 25327; Larry S. Nixon, 25640; Robert A. Vanderhye, 27076; James T. Hosmer, 30184; Robert W. Faris, 31352; Richard G. Basha, 22770; Mark E. Nusbaum, 32348; Michael J. Keenan, 32106; Bryan H. Davidson, 30251; Stanley C. Spooner, 27393; Leonard C. Mritchard, 29009; Duane M. Byers, 33363; Jeffry H. Nelson, 30481; John R. Lastovo, 33149; H. Warren Butman, Jr. 29366; Thomas E. Byrne, 32205; Mary J. Wilson, 32955; J. Scott Davidson, 33489; Alan M. Kagen, 36178; Robert A. Molan, 29834; B. J. Saddoff, 36663; James D. Berquist, 34776; Updeep S. Gill, 37334; Michael J. Shea, 34725; Donald L. Jackson, 41090; Michelle N. Lester, 32331; Frank P. Presta, 18828; Joseph S. Presta, 35329; Joseph A. Rhoa, 37515; Raymond Y. Mah, 41426. Also authorize Nixon & Vanderhye to delete any attorney names/numbers no longer with the firm and to act and rely solely on instructions directly communicated from the person, assignee, attorney, firm, or other organization sending instructions to Nixon & Vanderhye on behalf of the owner(s).

1. Inventor's Signature:	<u>Maurice</u>	<u>Peter</u>	<u>RUMSEWICZ</u>	Date: <u>15 Nov, 2000</u>
Inventor:	<u>Michael</u> (first)	<u>MI</u> (last)	<u>Australia</u>	<u>Stu X</u> Australian citizenship
Residence: (city)	<u>Victoria</u>			
Post Office Address:	<u>34 Remwick Street, Glen Iris, Victoria, Australia</u>			
(Zip Code)	<u>3146</u>			
2. Inventor's Signature:	<u>Maurice</u>	<u>Castro</u>	<u>CASSTRO</u>	Date: <u>16 /11 /2000</u>
Inventor:	<u>David</u> (first)	<u>MI</u> (last)	<u>Australia</u>	<u>Stu X</u> Australian citizenship
Residence: (city)	<u>Victoria</u>			
Post Office Address:	<u>3/26 Powell Street, South Yarra, Victoria, Australia</u>			
(Zip Code)	<u>3141</u>			

FOR ADDITIONAL INVENTORS, check box and attach sheet with same information and signature and date for each.